

Our Ref: Trans18

APPLICATION FOR UNITED STATES LETTERS PATENT

FOR

INTERPAGE PROLOGUE
TO PROTECT VIRTUAL ADDRESS MAPPINGS

Inventors:

Robert Bedichek
David Keppel
John Banning

Prepared by:

Stephen L. King
30 Sweetbay Road
Rancho Palos Verdes, California 90275
(310) 377-5073

Express Mail Label No. EK380384375US

INTERPAGE PROLOGUE TO PROTECT VIRTUAL ADDRESS MAPPINGS

BACKGROUND OF THE INVENTION

Field Of The Invention

5 This invention relates to computer systems and, more particularly, to methods and apparatus for assuring consistency of translated instructions being executed by a microprocessor which dynamically translates instructions from a target to a host instruction set.

History Of The Prior Art

10 Recently, a new microprocessor was developed which combines a simple but very fast host processor (called a "morph host") and software (referred to as "code morphing software") to execute application programs designed for a processor having an instruction set different than the instruction set of the morph host processor. The morph host processor
15 executes the code morphing software which translates the application programs dynamically into host processor instructions which are able to accomplish the purpose of the original software. As the instructions are translated, they are stored in a translation buffer where they may be executed without further translation. Although the initial translation of
20 a program is slow, once translated, many of the steps normally required for hardware to execute a program are eliminated. The new microprocessor has proven able to execute translated "target" programs as fast as the "target" processor for which the programs were designed.

The new microprocessor is described in detail in U. S. patent 5,832,205,
25 Memory Controller For A Microprocessor For Detecting A Failure Of

Speculation On The Physical Nature Of A Component Being Addressed,

Kelly et al, November 3, 1998, assigned to the assignee of the present invention.

One reason that the new processor is able to execute programs rapidly is
5 its ability to link together sequences of translations that occur frequently
into very long sequences. Linking eliminates many of the steps which
would be necessary to retrieve individually the various translations for
execution. The process by which this is accomplished is explained in
detail in the above-mentioned patent.

10 One problem that must be resolved for a computer which executes host
translations of a target program is that the target program typically
defines the sequences of target instructions which are to be executed by
presenting a series of addresses at which those instructions are stored to
the central processor as those target instructions are to be executed.
15 The central processor reads the address of the instruction next to be
executed, fetches that instruction from memory, and executes the
instruction. When the target program being executed is defined by such
a sequence of addresses yet the instructions being executed are host
translations of those instructions which reside at other addresses, it is
20 necessary to determine that each translated host instruction is, in fact,
the result of a translation from a target instruction which is at the
address (including the effect of address mapping) presented by the target
program for execution.

This is an especially difficult problem where sequences of translated instruction have been linked together in the manner described above in order to attain rapid execution.

It is desirable to improve the operation of a computer system which
5 utilizes a microprocessor that translates programs dynamically from target instructions into host instructions able to accomplish the purpose of the original software by rapidly determining that a host instruction is a translation of a target instruction presented for execution.

Summary Of The Invention

10 It is, therefore, an object of the present invention to improve the operation of a computer system which utilizes a microprocessor to translate programs dynamically from target instructions into host microprocessor instructions able to accomplish the purpose of the original software by rapidly determining that a host instruction is a
15 translation of a target instruction presented for execution.

This and other objects of the present invention are realized in a computer which translates instructions from a target instruction set to a host instruction set by a process for testing the memory address of a target instruction to be executed against a copy of the memory address of the
20 target instruction from which a translation of the target instruction was made, executing the translation if the addresses compare, and generating an exception if the addresses do not compare.

These and other objects and features of the invention will be better understood by reference to the detailed description which follows taken

together with the drawings in which like elements are referred to by like designations throughout the several views.

Brief Description Of The Drawings

Figures 1 is a block diagram illustrating a new microprocessor which is adapted to carry out the present invention.

Figure 2 is a diagram illustrating some of the steps of a process for carrying out the invention.

Figure 3 is a diagram illustrating some of the steps of another process for carrying out the invention.

Figure 4 is another diagram illustrating the steps of a general process for carrying out the invention.

Detailed Description

Figure 1 illustrates a microprocessor 11 which may utilize the present invention. The microprocessor pictured is described in detail in U. S. patent 5,832,205. The microprocessor includes a floating point unit 12, an integer unit 13, a translation buffer 14 which is a part of system memory, target memory 15 which is another portion of system memory, and a translation lookaside buffer 18 which is a part of a memory management unit 19.

As described above, the new microprocessor translates target instructions into host instructions which it executes. The target instructions are stored at addresses in a target portion of system memory while translated host instructions are stored at different addresses in a

host portion of system memory. When a target program is being executed, an instruction pointer pointing to the address of each sequential target instruction to be executed is provided in an EIP register. The new processor reads the instruction pointer, determines
5 the address of the target instruction, and decides whether a host translation exists for that instruction. If no translation exists, then a new translation is generated, stored in host memory, and executed.

If a translation exists, the code morphing software finds the translation and executes it. If the translation is not linked to any other translations,
10 then after it has been executed software determines the next target instruction from the next instruction pointer (often herein referred to as the EIP value) in the EIP register, determines whether a translation exists for this next target instruction, and continues with the process in the manner just described.

On the other hand, if the translation which is being executed is linked to
15 other translations, then the next host instruction to be executed is not determined by the target instruction EIP value. Instead, the next host instruction is determined by a jump command to the next translation placed at the end of the executing translation by a linking process of the
20 code morphing software.

In each case in which a next host translation is to be executed, it is important for the software to ascertain that the translation which is to be executed is a translation of the target instruction commanded by the target program which is being executed by means of the host translation
25 and that both the target instruction EIP value and the physical address it

maps to correspond and are correct. The translation to be executed must be a translation of the target instruction to which the EIP value points and which is stored at the physical address the EIP maps to if execution is to produce a correct result.

5 In order to be able to ascertain that the translation which is to be executed is a translation of the target instruction pointed to by the EIP value and stored at the physical address the EIP value maps to, the present invention stores for each translation the physical address in memory of the target instruction from which the translation was made,
10 the EIP of that target instruction, and various context relating to the translation not pertinent to the present invention.

When the instruction to be executed is being "dispatched" (responding to the next target instruction of the target program), the operation proceeds in a straightforward manner as shown in Figure 2. The code morphing
15 software simply maps the EIP value to a physical address and finds the translation with the same EIP value and physical address. To accomplish this, the value in the EIP register is furnished to the code morphing software. The code morphing software adds the code segment base and the EIP value and furnishes the result to the translation
20 lookaside buffer circuitry. The translation lookaside buffer circuitry utilizes this logical address to determine the physical page address pointed to by the EIP value. The physical page address and the low order bits of the sum of the code segment base and the EIP value are then combined to provide the physical address of the target instruction. With
25 this physical address, a translation is found that matches this physical address and EIP value. If such a translation is found, control is

transferred to it. If none is found, a new translation is made of the target instructions starting at that EIP value and physical address.

However, when a first host translation is linked to one or more succeeding translations, the operation is much more complicated.

5 Although the EIP value of the first translation in the sequence will have been used to check that its physical address and the physical address stored with the first host translation are in fact the same, the next translation (translated instruction) is reached by a linking jump instruction which completes the execution of the first translation. Since
10 such a jump instruction does not return control to the dispatching portion of the program but rather intentionally eliminates that portion of the program in order to attain speed, no check would normally be made to determine that the host instruction is a translation of the target instruction stored at the memory address pointed to by the EIP value at
15 the physical address mapped to by the EIP value.

Because the EIP value is changed in a fixed and predictable manner, it can be guaranteed that the EIP value of the second translation is correct when the two translations are chained together. However, the mapping of the EIP value of the second translation to a physical address may not
20 be the same as the mapping of the EIP value of the first translation to a physical address so something must be done to assure that the mapping of the EIP value of the second translation is to the physical address of the second translation.

The present invention assures that such check is, in fact, carried out by
25 providing a process for accomplishing the consistency test in all

situations where required. The jump instruction which accomplishes the linking is caused to point to a process which accomplishes the check and then moves to the linked translation if the test is met. If the test is not met, an exception is generated; and various solutions may be provided.

5 It should be noted that the process of the present invention also provides a solution to situations in which an operating system has for some reason unmapped a particular memory page at which translated target instruction are stored. The prologue process will indicate a failure when its test is conducted so that the link to the next translation will not be
10 executed. At that point, the code morphing software may roll back execution to the last consistent point from which some other method for executing the target instructions may take place.

In one embodiment of the invention, the process (shown in Figure 3) for accomplishing the check utilizes the EIP value designating the next
15 target instruction and the code segment base to determine the virtual address pointed to by the EIP value, determines the physical address of the target instruction by a lookup, and then compares this physical address with the physical address of the target instruction stored for the translation. If the addresses are the same, the linked translation is
20 executed. If the addresses differ, an exception is generated; and execution is rolled back to the last point at which consistent state existed and some other method of executing the target instruction is undertaken. At that point, the succeeding translation can be invalidated; and another translation or an interpretation of the target
25 instructions can be undertaken. Alternatively, the succeeding translation may remain valid, but its link to the preceding translation

broken so that the first translation no longer proceeds directly to the second.

The need for this general test may be eliminated if the EIP value for the succeeding translation maps to an address on the same memory page as the EIP stored for the preceding translation because the consistency test will already have been conducted for the memory page and the prior translation will not have changed the mapping. If the logical addresses are on the same memory page, then the physical addresses are also on the same memory page. The information to determine this is available from the EIP values of the preceding and succeeding translations and the physical address of the first target instruction of the preceding translations. The EIP value furnished by the target program and the physical address provide sufficient information to determine whether the stored EIP values saved for the host translation are on the same memory page.

The process of the present invention may be utilized in at least two distinct ways which are both shown in figure 3 of the drawing. If it is known that the translation will be linked to a previous translation when it is translated and a check of the physical address is required, then the prologue process (described above and illustrated in Figure 3) may be included in the translation. In such as case, the preceding translation merely jumps to the next translation where the address consistency process of the prologue is executed before the translation is executed.

If, on the other hand, it is not known when the translation is made whether it will be linked to another translation or not, then the prologue

process is generated as a separate short prologue when the linking of the two translations occurs. At this time, the preceding translation is provided a jump instruction to the prologue process; and the prologue completes (if the test is met) with another jump instruction to the succeeding translation.

Another manner of utilizing the prologue process would be to incorporate what started as a separate prologue into a linked translation at the time of linking.

A single translation may include translations from target code which appears on two different pages. In this case, the translation must include code for performing the EIP mapping and physical address consistency check described above for target instructions on any page other than the page of the first target instruction.

Although the present invention has been described in terms of a preferred embodiment, it will be appreciated that various modifications and alterations might be made by those skilled in the art without departing from the spirit and scope of the invention. The invention should therefore be measured in terms of the claims which follow.

What Is Claimed Is: